Both Questionable and Open Research Practices are Prevalent in Education Research

Matthew C. Makel <https://orcid.org/0000-0002-3837-0088>

Jaret Hodges <https://orcid.org/0000-0002-4427-5247>

Bryan G. Cook <https://orcid.org/0000-0001-9294-0873>

Jonathan A. Plucker <https://orcid.org/0000-0002-5327-0851>

This is a preprint.

The version of record is published in *Educational Researcher* and can be found at:

<https://doi.org/10.3102/0013189X211001356>

Abstract

Concerns about the conduct of research are pervasive in many fields, including education. In this preregistered study, we replicated and extended previous studies from other fields by asking education researchers about 10 questionable research practices and 5 open research practices. We asked them to estimate the prevalence of the practices in the field, to self-report their own use of such practices, and to estimate the appropriateness of these behaviors in education research. We made predictions under four umbrella categories: comparison to psychology, geographic location, career stage, and quantitative orientation. Broadly, our results suggest that both questionable and open research practices are used by many educational researchers. This baseline information will be useful as education researchers seek to understand existing social norms and grapple with whether and how to improve research practices.

Keywords: Open Science, Questionable Research Practices, Replication

Author Note: Preregistration, survey materials, code, and data can be found at:

<https://osf.io/83mwk/>

Acknowledgement: We thank Dr. Hannah Fraser and Dr. Fiona Fidler for sharing their survey and code and Dr. Leslie John for sharing data and code.

Author contact: makel@jhu.edu

Both Questionable and Open Research Practices are Prevalent in Education Research

In recent years, education researchers have reflected on their processes for conducting research and disseminating results (e.g., Pigott et al., 2013). This self-reflection is more advanced in other fields. For example, in psychology, this introspection catalyzed what has been called a replication crisis (Pashler & Harris, 2012, p. 531) due to concerns that many seemingly well-established findings could not be independently replicated. Perhaps the most telling example of the lack of replicability is the Open Science Collaboration's Reproducibility Project: Psychology (OSC, 2015), in which 270 researchers collaborated to replicate 100 studies published in major psychology journals. The OSC preregistered their protocols and used sample sizes larger than the original studies. Only 39% of the replication attempts found results similar to the original studies. Moreover, the effect sizes in the replication attempts were typically less than half the size of the original results (see also Camerer et al., 2018; Klein et al., 2018). These results suggest neither the pattern nor magnitude of many published results in psychology are replicable.

Psychology is not alone. Researchers in preclinical cancer trials and drug trials have had similarly unimpressive results when trying to replicate previously published findings (Begley & Ellis, 2012; Prinz et al., 2011). If neither the pattern nor magnitude of many published research findings is reliable, that should be of great concern to researchers who use those findings in their own work, funders who want to support effective practice, and, perhaps most importantly, practitioners and policymakers who rely on that research to guide application. Replicability of research findings is especially important in applied fields such as education, where research often serves as the foundation for policy and practice. In this paper, we explore several facets related to whether education research may have similar problems.

Why are published research findings so unreliable? Many reasons have been proposed, including publication bias (Polanin et al., 2016; Sham & Smith, 2014), selective reporting of results (Pigott et al., 2013), and the file drawer problem of not submitting null findings (Rosenthal, 1979; Rotton et al., 1995; Spellman, 2012). Bias against null or negative results is so strong that they appear to be “disappearing from most disciplines and countries” (Fanelli, 2011, p. 891). Another relevant concern is the quality of the methods used to produce initial findings. Simmons et al. (2011) demonstrated that researchers have the flexibility to manipulate their data, which can lead to even highly improbable, desired results (in their case, that listening to a Beatles song makes the listener younger). This flexibility includes how much data to collect, what data and observations to exclude, and how to analyze data and report findings. They later wrote that such practices were wrong, but not wrong in the way jaywalking is wrong, wrong in the way that robbing a bank is wrong (Simmons et al., 2018).

Questionable Research Practices

Although Simmons and colleagues (2011) demonstrated that such manipulations are easy to engage in, they did not establish the prevalence of such undesirable practices. To examine prevalence, John and colleagues (2012) administered a survey to several thousand psychologists at U.S. universities asking them whether they had ever engaged in a broad spectrum of research behaviors deemed questionable research practices (QRPs). They are questionable because they fall into a “gray area of acceptable practice” with no consensus on whether they are acceptable (p. 524; see also Butler et al., 2017). Such behaviors included collecting more data after checking to see if the results were (already) significant, excluding data after checking its impact on results, and hypothesizing after the results are known (HARKing; Kerr, 1998). The self-admission rate for QRPs ranged from 1% for falsifying data to 63% for failing to report all dependent measures

used in a study, with more than 25% of participants reporting they had engaged in at least 7 of the 10 practices. Estimates of QRP prevalence have since been replicated with samples of German (Fiedler & Schwarz, 2016) and Italian psychologists (Agnoli et al., 2017). They have also been collected in ecology (Fraser et al., 2018). Although each study's methods and results varied slightly from the others, the general trend of researcher self-admission of QRPs was consistent across all studies.¹

Open-Science Practices

In response to low replication rates, QRPs, and related issues, many research communities have begun implementing a series of open-science reforms (e.g., Nosek et al., 2015). Broadly, open science efforts seek to make the scientific research process more accessible and transparent through practices such as preregistering predictions and analysis plans, sharing research materials such as code and survey items, sharing data, and making research findings publicly available. By making science open, the hope is to align scientific behaviors and incentives with scientific values to foster a “credibility revolution” in scientific research (Vazire, 2018, p. 411).

This methodological self-reflection has led to changes such as requiring all clinical trials to preregister their goals (Kaplan & Irvin, 2015) and the creation of high-profile, open-access journals such as *AERA Open* and *PLoS*, as well as the proliferation of preprint and postprint servers that remove paywall barriers to increase accessibility to academic research (e.g., EdArXiv.org). Similarly, there has been a call to shift emphasis away from using statistical significance as a criterion for publication, to assessing the quality of the research question and the methods used to answer it. Registered reports are one method of accomplishing this

¹Whether any QRPs are a form of scientific misconduct is a question beyond the scope of this manuscript, but we acknowledge that this is an important ethical debate.

(Chambers et al., 2014; Gehlbach & Robinson, 2018). As of January 2021, over 275 journals accept the registered report manuscript format, in which peer review occurs prior to data collection. Within education, *AERA Open* and *Journal of Computer Assisted Learning* have run special topic issues with registered reports, and *Journal of Educational Psychology*, *Exceptional Children*, *Gifted Child Quarterly*, and the *British Journal of Educational Psychology* accept them as regular submissions. Although benefits and obstacles to adopting open-science practices exist for researchers at all career stages, early-career researchers may be particularly receptive to engaging in open science (Allen & Mehler, 2019; Farnham et al., 2017). Indeed, research has found early-career researchers are more likely to engage in open-science practices such as data sharing (Campbell et al., 2019).

Calls for greater use of open-science practices to address these issues have already begun in education in general (van der Zee & Reich, 2018), special education (Cook et al., 2018), and gifted education (McBee et al., 2018). Currently, we do not know how many education researchers use open practices or QRPs. Nor do we know whether such practices are viewed by the education research community as typical and/or appropriate.

Relevance to Education Research

Although education research has its own history, community, and standards of practice, it is not immune to the aforementioned problems in other disciplines. For example, to estimate how often replications were conducted in education research, Makel and Plucker (2014) reviewed the top 100 education journals as ranked by their impact factor. Only 0.13% of articles explicitly claimed to be replications. They concluded that conducting more replications, especially independent replications, will help education research serve as a reliable resource for developing sound policy and practice. Additionally, Pigott et al. (2013) reported that non-significant

outcomes in education dissertations were 30% more likely than significant findings to be omitted from published articles (i.e., outcomes-reporting bias). And Polanin et al. (2016) found that effect sizes for unpublished studies included in education meta-analysis were significantly smaller than for published studies, suggesting the research base in education may reflect publication bias in favor of large, positive findings. These examples demonstrate that the published record in education research should be consumed with caution. QRPs may be an invisible thread weaving these examples together. However, the lack of transparency around most education publications (e.g., were all comparisons reported?) inhibits understanding of the presence of QRPs and their impact on the published literature. Similarly, little evidence exists about the prevalence of open-science practices that might curtail the use of QRPs in education.

Present Study

Despite the attention QRPs have received in other academic communities and the calls for greater use of open-science practices in education, researchers have not examined how often such behaviors occur in education research. In the present study, we sought to replicate and extend previous QRP studies. Using Fraser et al.'s (2018) QRP items along with five new items on open-science practices, we asked participants to (a) estimate the percentage of researchers in education who have engaged in each research practice, (b) estimate how often *they* have engaged in each practice themselves, (c) share their opinion of whether each practice should be used, and (d) describe why they think each practice should or should not be used. Answering these questions (a) creates a rough estimate of how many education researchers have used QRPs and Open Practices, (b) estimates current social norms for the field on how often these practices are used, and (c) estimates perceptions of how often it is acceptable to use these practices. If

researchers believe a behavior is a social norm (regardless of whether the belief is accurate), then that belief may influence their likelihood of engaging in the behavior.

Limitations

At the request of the editor, we discuss limitations before describing study methods to highlight several limitations that likely contribute to making our results relatively noisy estimates of actual beliefs and behaviors in the education research community. These limitations include low response rate, measurement error, potential time confounds, and need for independent verification.

Our approximately 10% response rate raises reasonable concerns about representativeness and generalizability. It could be that the most frequent users of QRPs and open practices ignored our survey, thus making our results underestimates of prevalence rates of both. It could also be that only the most frequent users of QRP and open practices responded, making our results an overestimation of these practices in education research. Which end of the spectrum our results fall closer to is unknown, although previous research on self-reports of potentially unethical and ethical behavior and our bootstrapping estimates (see online supplement) give us confidence that the results fall somewhere in the middle, with a likelihood of an underestimation of QRP use and an overestimation of open practice use (see Choshen-Hillel et al., 2020; Halevy et al., 2014). Additional surveys, focus groups, and analysis of specific publications or bodies of research will help better inform this understanding and assess the generalizability and replicability of our findings. It is worth noting that the issue of response rates also came up in John et al. (2012). They wrote in their online supplement that “some have argued that response rates may not play as important a role in determining the validity of

prevalence estimates as participants' perceptions of anonymity and confidentiality (Bates & Cox, 2008; Fanelli, 2009)" (p. 6).

As suggested by a reviewer, we do not know whether our survey captures accurate estimates of past behavior or is "capturing some underlying dimension of people's memory of the frequency with which they do things" (Reviewer 1). Method of inquiry, sampling bias, and other measurement-related factors may influence results. Assessment of such validity concerns requires additional research and external validation connecting self-reports and actual behavior. Our belief is that given existing social biases, participants were likely to underreport rather than overreport QRPs. If we accept Simmons et al.'s (2018) jaywalking-versus-bank robbing analogy as relevant, we suspect most individuals would accurately recall whether they had ever robbed a bank. We also suspect most individuals would accurately recall whether they had ever jaywalked.

Our current data are confounded by time factors. We assessed current beliefs regarding appropriateness but asked respondents to report whether they had ever used the practices. Our data do not tell us the person's beliefs at the time they engaged in the practice. Such work would help reveal whether researchers have changed their beliefs or whether their beliefs and actions did not align at the time.

Finally, our results do not provide an estimate of how many research findings are the result of QRPs and should not be interpreted as such. Admitting to using QRPs does not automatically invalidate all of one's research (perhaps not even for any specific study). Furthermore, we do not know the extent to which social desirability bias may affect self-report rates and we did not assess self-reported motivations for QRP use. There may be instances where QRP use was appropriate or encouraged externally (e.g., reviewers or editors). That said, if

consumers of research are making decisions about appropriateness of practices that may be questionable, what was done must be communicated transparently. Open practices facilitate such transparency.

Method

We emailed potential participants an anonymous link to a survey in Qualtrics asking about common statistical analysis and reporting practices in education research. The email stated that we were interested in opinions only and that it was not a statistics test. Two follow-up email reminders were sent. No compensation was provided. Participants were given an option to receive a summary of results.

Sample

We emailed researchers who had published at least one article within the last 10 years (2008-2018) in a journal published by the American Educational Research Association (AERA) or an academic education research journal that was ranked in the top 30 by impact factor as of 2017 by Scimago Lab. This group was selected to represent typical researchers active within the education research community. Additional sample details and journal inclusion/exclusion criteria are available in the online supplement and Tables S1 and S2, respectively.

All authors with contact information associated with the article were emailed (i.e., if the journal listed contact information for a single author, then only that author was contacted; if multiple authors' contact information was provided, multiple authors were contacted). We sent messages to 14,228 email addresses. Of those, 814 emails failed or bounced, and over 2,700 resulted in an automatic email reply that included failed delivery or some form of "out of office" (vacation, family leave, sabbatical, moved institutions) message. After updating as many email addresses as feasible, we ended up with responses from 1,488 participants (see online

supplement for details). As explained below, many did not complete the entire survey, but we obtained roughly the same number of responses across all items because survey item order was randomized.

Prior to data collection, we conducted a power analysis. The power curve for hypothetical sample sizes can be found in Figure S1 and shows that a sample of over 1,000 would have sufficient power to identify small effect sizes.

Materials

After a general introduction of the survey, we asked participants for consent and whether they had participated in conducting quantitative research that involved null-hypothesis significance testing (i.e., calculating p -values). The response options for the latter question were “Yes” and “No, my research is qualitative, descriptive, single-case design, case studies, and/or theoretical or position-pieces.” Those who answered yes were considered quantitative researchers, those who answered no were considered non-quantitative researchers. Due to the potentially sensitive nature of their responses and to maintain confidence in anonymity, we asked participants only two demographic questions at the end of the survey: (a) their broad career stage and (b) whether they were based professionally within the United States.

We used 15 item stems to ask four repeated questions. For QRPs, we used the first 10 items from Fraser et al. (2018). For Open Practices, we created five new items (preregistration, open data, open materials, replication, and open access) specifically for this study (see Table 1 for exact wording of items). Quantitative researchers received all 15 items; non-quantitative researchers received all items that were not explicitly about p -values (omitting QRP items 1, 2, 5, 6, 7, and 8). The item stems were presented in randomized order across participants.

The four repeated questions asked of each item stem were prevalence estimates (the percent of researchers using this practice), self-reported use (Never = 0, Once = 1, Occasionally = 2, Frequently = 3, and Almost Always = 4), appropriateness estimates (whether the practice should ever be used; 0 = Never, 1 = Rarely, 2 = Often, 3 = Almost Always), and an open-ended question on why the practice should be used. We did not analyze the open-ended responses; a subsequent manuscript will report those results.

Hypotheses

Much of this study is exploratory and descriptive in nature. To our knowledge, there is no relevant theory and only a few related studies from other fields. As such, many of our predictions are “hunch based” and not theory based. We include them to differentiate our *a priori* hunches from any developed after seeing the data. Specific confirmatory predictions about each repeated question were preregistered (<https://osf.io/83mwk/>) and vary across four categories: comparison to psychology; geographic location; career stage; and quantitative orientation. We outline our predictions here but see the Online Supplement for a complete reporting.

Comparison to psychology. Because John et al. (2012) reported QRP rates as yes/no, we could only compare rates on a dichotomous basis. We predicted equivalence between psychology (John et al., 2012, Table S2, control group) and education researchers in prevalence estimates for QRPs within their fields and for self-report. We had no *a priori* predictions for appropriateness estimates.

Geographic location. Because Fanelli et al. (2015) found variation in scientific misconduct rates based on geographic location but others did not find such variation in QRP use (e.g., Agnoli et al., 2017), we wanted to assess whether there was variation based on geographic location. When comparing U.S. education researchers versus those based outside the U.S., we

predicted equivalence on prevalence estimates of QRP use in the field, but that U.S.-based researchers would estimate higher use of Open Practices. We also predicted that U.S.-based researchers would self-report lower QRP use and self-report higher use of Open Practices. We had no *a priori* predictions about appropriateness estimates for QRP or Open Practices.

Career stage. Similarly, because variations in scientific misconduct (Fanelli et al., 2015) and adoption of Open Practices (Campbell et al., 2019) have been reported based on career stage, we wanted to assess whether there was variation based on career stage. We predicted higher prevalence estimates of QRPs and Open Practices at the early career stage. In contrast, we predicted lower self-reports of QRPs, but higher self-reports of Open Practices at earlier career stages. However, we predicted that appropriateness estimates of QRPs would be more positive in later career stages, but Open Practices would be higher at earlier career stages.

Quantitative orientation. Both quantitative and non-quantitative research is common within education and each has its established behavioral norms. We wanted to assess whether there was variation between the two in QRP and Open Practices. We made no specific confirmatory hypotheses about the responses of quantitative versus non-quantitative researchers. All comparisons are exploratory; no p-value comparisons were made, but effect sizes were calculated comparing responses of those who self-identified as having conducted quantitative research with those who self-identified as having done non-quantitative research.

Analysis

We preregistered that we would use ANOVAs to test our hypotheses of differences, and equivalence tests when we predicted equivalence. A two-category ANOVA was used for geographic comparisons as well as to compare the responses of self-identified quantitative and non-quantitative researchers. A three-category ANOVA was used for career stage analyses.

Whereas traditional frequentist statistics report on the evidence of *difference*, equivalence tests assess the evidence of *similarity*. The equivalence test (Lakens, 2017) can be thought of as an inverted t-test. An important consideration in testing for equivalence is determining equivalence bounds. Lakens suggested determining bounds through either an expected effect size difference or through a power analysis. We used the first approach because our expected sample size ($n > 1,400$) would make virtually any difference statistically significant but not necessarily practically significant. Using Lakens' approach, we set equivalence bounds at an effect size of $d = 0.10$. Lakens referred to this estimate as the smallest effect size of interest. In other words, our test designated results as equivalent when the mean-difference effect size was 0.10 or less, representing a frequency difference of 5% (e.g., if the difference in the frequency of a QRP between two groups is greater than 5%, an equivalence test provides evidence that those two groups are not equivalent).

Effect sizes. We calculated effect size estimates for our analyses around comparisons to psychology, geographic location, career stage, and quantitative orientation. For comparison of two groups (i.e., comparisons to psychology, geographic location, and quantitative orientation) we calculated Cohen's d ; for comparisons of three groups (i.e., career stages), we calculated η^2 (Lakens, 2013).

Results

We first report overall descriptive results. Next, we report results based on our predictions under each of the four umbrella categories (comparison to psychology, geographic location, career stage, and quantitative orientation). Full results can be found in Table 1 and the Online Supplement.

Overall Descriptive Results

Prevalence estimates. Prevalence estimates demonstrate the perceptions of how typical a practice is within the field; practices with higher prevalence estimates are likely perceived as normative, whereas practices with lower prevalence estimates are likely viewed as atypical. As shown in Figure 1, prevalence estimates of QRPs ranged from 23% (Filling in Missing Data) to 61% (Omitting Nonsignificant Studies or Variables). These results suggest several QRPs are believed to be fairly typical within education research. Prevalence estimates of Open Practices ranged from 25% (Data Sharing) to 50% (Open Access). These results provide evidence that Open Practices are perceived to be used by a non-trivial minority of education researchers by our sample. Our results do not reveal whether prevalence estimates are accurate, they only establish what is perceived as normal behavior. However, perceptions of others' behavior may influence researchers' attitudes and behavior, and thus may be important regardless of their accuracy. Regardless, increasing awareness of the prevalence estimates may help change subsequent behavior.

Self-reported use. As shown in Figure 2, self-reports of using QRPs at least one time ranged from 10% (Filling in Missing Data) to 67% (Omitting Analyses). These results suggest that QRPs are fairly common within our sample, although reported use varies substantially across practices. Self-reports of using Open Practices at least one time ranged from 43% (Replication) to 78% (Open Access), suggesting that Open Practices are common within our sample, although reported use varies substantially across practices. These results help lay the groundwork for knowing prevalence of behaviors in education research. If QRP use is this common, then policy, practice, and research that are based on education research findings may not be as strongly empirically supported as consumers believed. This could meaningfully impact expected results of interventions, effectiveness of policy, and likelihood of replicability.

An estimate of how many participants reported never using any QRP could be either 4.4% or 8.6% depending on whether we use all responses or only those who responded to all items, respectively. On the other hand, an estimate of how many participants reported never using any Open Practice could be either 3.8% or 6.8% depending on whether we used all responses or only those who responded to all items, respectively. These results suggest that nearly all education researchers in our sample have used QRPs and Open Practices. These estimates do not establish whether the behaviors are appropriate.

Appropriateness estimates. Appropriateness estimates demonstrate the perceptions of whether a practice should be used within the field. We report the percent of participants responding that a practice should never be used (higher rates mean less appropriate). As shown in Figure 3, the proportion of respondents indicating that QRPs should never be used ranged from 26% (Omitting Analyses) to 82% (Filling in Missing Data). These results indicate great variability in the perceived appropriateness of various QRPs, with many believing there can be instances of appropriate use. These varied results suggest there is no clear consensus of whether most QRPs are appropriate or misconduct; they remain questionable. Participants indicating that an Open Practice should never be used ranged from 2% (Replication) to 5% (Open Access). These results suggest that participants generally viewed Open Practices as appropriate for education research. Given this context, if greater use of Open Practices is desired, future intervention to make them easier to implement and more rewarded may be more fruitful than efforts to persuade education researchers of their appropriateness, as the researchers in our sample appear to need little persuasion. Together, the QRP and Open Practices appropriateness results are relevant because they establish a foundation for education researchers' beliefs about

what behaviors are social norms for the field. For example, is each QRP misconduct (never should be done), a best practice (done as often as possible), or something in between?

Summary of Predicted Comparisons

As described more fully in the Online Supplement (Tables S6-S10), we made predictions about prevalence estimates, self-reported use, and appropriateness estimates of QRP and Open Practice use under four umbrella categories: comparison to psychology, geographic location, career stage, and quantitative orientation. Here we provide a summary of those results.

Comparison to psychology. As predicted, prevalence estimates and self-reported use of QRPs were largely equivalent between our sample and John et al.'s (2012) sample of psychologists. There were a few notable exceptions, such as our participants admitted to Filling in Missing Data at higher rates (10%) than psychology participants (0.61% [$d = 0.73$]). For appropriateness estimates, we made no predictions, but there were practical differences ($> 5\%$) across all items. Our participants were more likely than psychologists to report that a QRP should never be used, with the exception of Filling in Missing Data. Caution should be used in interpreting these results. John et al. (2012) asked whether a given practice was defensible. In comparison, using Fraser et al.'s (2018) wording, our survey asked for participants' opinions on how often a given practice should be used. The identified differences could stem from the wording discrepancy, not a difference in beliefs across education and psychology researchers.

These results indicate that QRPs are reported at similar rates in education and psychology, a field with a recognized replication crisis (Pashler & Harris, 2012). Moreover, our results suggest that education researchers may feel that QRPs are less appropriate than did psychology researchers in 2012, which may suggest that the education research community will be relatively receptive to reforms aimed at curbing the use of QRPs.

Geographic location. Our prediction of equivalence for prevalence estimates of QRPs based on whether researchers were based inside the U.S. was generally supported (see Table S6 for full results); with only Data Peeking not equivalent (34% and 40% for U.S. and non-U.S. participants, respectively [$d = -0.24$]). Similarly, prevalence estimates of four Open Practices were not statistically significantly different from each other (contrary to our prediction); with only Preregistration showing statistical significance (27% and 34% for U.S. and non-U.S. participants, respectively [$d = -0.26$]), but in the direction opposite than predicted. We predicted self-reported use of QRPs would be lower and Open practices higher in the U.S. than outside the U.S. Self-reported use of six of the 10 QRPs matched this prediction, whereas four of the five Open Practices showed no statistically significant differences based on geographic location. The fifth, Preregistration (U.S. = 1.19, Non-U.S. = 1.54, $p > .999$ [$d = -0.24$]), appeared to have a difference, but in the opposite direction of our prediction (lower in U.S.). We made no predictions about appropriateness estimate differences based on geographic location. Although our exploratory analyses revealed many potential differences, no clear pattern emerged. Although cultural patterns may exist, our methods did not reveal them. Greater understanding of cultural patterns can help inform causes of researchers' behaviors as well as potential targeted interventions to change researcher behavior.

Career stage. Our prediction that earlier career researchers would estimate both QRPs and Open Practices as more prevalent than their more experienced peers received a small amount of support. When statistically significant differences were observed (only three of the 10 QRPs), earlier career stage participants reported higher prevalence estimates of QRP than later career stage participants (see Tables S7 and S8 for full results). Contrary to our predictions, participants' prevalence estimates of open practices did not differ across career stage. Contrary to

our predictions for self-reported use, we did not find statistically significant differences on self-reported use of most QRPs across career stage. Similarly, early career participants reported lower rates of using Open Practices than later career participants, contrary to our predictions. Further, contrary to our prediction, appropriateness estimates of QRPs and Open Practices were generally not significantly different from each other based on career stage. Our data did not reveal clear pattern of career stage being associated with beliefs and behaviors. Again, such patterns may exist, but our methods did not reveal them. Additional understanding of whether and how perceptions and behaviors regarding QRPs and Open Practices vary by career stage for education researchers can help inform efforts to change those perceptions and behaviors.

Quantitative orientation. We made no predictions about differences between quantitative and non-quantitative researchers (see Table S9 for full results). Nearly all prevalence-estimate and self-report gaps for both QRPs and Open Practices were smaller than our practical significance threshold. Nearly all appropriateness-estimate gaps for QRPs were smaller than our practical threshold. All reported appropriateness estimates for Open Practices exceeded our set practical threshold between quantitative and non-quantitative researchers, with quantitative researchers rating each Open Practice as more appropriate than non-quantitative researchers. Differences in perceptions across research orientation may represent appropriate distinctions for each area or may indicate differing views or goals in research (e.g., prediction vs. explanation). These findings are important because they suggest researchers' methodological orientation may be a relevant variable to consider when examining perceptions and use of Open Practices.

Discussion

Other fields have established that QRPs are not uncommon (e.g., Fraser et al., 2018; John et al., 2012), and open-science practices have been proposed as an antidote to increase transparency (Nosek et al., 2015). We asked education researchers to estimate the percentage of their colleagues who used questionable and open research practices (prevalence estimate), how often they used them themselves (self-report), and how often such practices should be used (appropriateness estimate). Self-reported use of QRPs varied, with only one under 10% (Filling in Missing Data) to as high as 67% (Omitting Analyses). Self-reported use of Open Practices varied from 43% (Replications) to 78% (Open Access). Collectively, these results suggest both QRPs and Open Practices have been used by many education researchers in our sample.

Nonetheless, estimates of prevalence within the field of QRPs and Open Practices varied. For QRPs, estimates ranged from 23-61% (Filling in Missing Data and Omitting Nonsignificant Studies or Variables, respectively). For Open Practices, prevalence estimates ranged from 25-50% (Data Sharing and Open Access, respectively). Our data do not reveal whether these self-reported estimates are accurate. Regardless of their accuracy, findings suggest that participants believe that both QRPs and Open Practices are relatively common in the field. If a behavior is perceived as typical, there may be less concern about the stigma of violating a social norm. Such perceptions may affect adoption of research practices by individual researchers.

Although we did not plan formal comparisons across prevalence estimates and self-report use, prevalence estimates generally appear higher than self-reported use for QRPs, but lower for Open Practices. Such gaps indicate that participants tend to believe their peers are more likely to use QRPs and less likely to use Open Practices than themselves. Thus, participants may have an artificially bleak picture of the use of QRPs and Open Practices in the field that could serve to make QRPs more socially acceptable and Open Practices less socially acceptable because of

(potentially inaccurate) perceptions about normative behavior. Alternatively, as investigated by John et al. (2012) in psychology, these apparent gaps could demonstrate a desirability response bias despite the survey being anonymous.

The percentage of participants indicating that a practice should never be used varied considerably across QRPs (26% for Omitting Analyses up to 82% for Filling in Missing Data). Very few participants reported that Open Practices should never be used (2% for Sharing Materials up to 5% for Open Access). The QRP results suggest that, like in psychology, these practices fall in a gray area for the education research community, with many individuals believing they are inappropriate and others seeing situations in which they may be appropriate to use. Such disagreement is to be expected but given that application of QRPs can lead to accuracy, generalizability, and replicability problems, careful consideration of how the field should address the situation is warranted. If the response is to do nothing, errors are likely to remain and the source of some of those errors will remain hidden.

These appropriateness estimates serve as a form of assessment of the perceived gravity of each practice. Simmons et al. (2018) used to view them as equivalent to jaywalking (i.e., against the rules, but relatively harmless), but then changed their minds to view many QRPs as equivalent to robbing a bank (against the rules and harmful to others). Education researchers must wrestle with this jaywalking-versus-bank robbing analogy. Should school districts adopt an intervention to close an achievement gap if researchers had conducted 20 analyses on the data, but only reported the results that make the intervention look effective? We do not believe so. But given the self-reported use of QRP responses, we suspect many such scenarios occur.

One reviewer commented, “the replicability problem is likely to be driven more by unknowingly shoddy work than outright subversiveness.” Our current data do not address

whether research behaviors occur due to ignorance or purposeful deviation from norms. But if the field wishes to change behaviors (or beliefs), knowing the rationale behind research behaviors could help create more effective course of action. Such actions could include training and professional development; changes to incentives; or changes to processes like requiring more preregistration, incentivizing registered reports, or requiring disclosure regarding specific questions about manuscripts (e.g., are all analyses reported?).

The Open Practices self-report findings are encouraging. However, the findings related to replication research, for example, are inconsistent with reports that only 0.1% of journal articles in education journals are replications studies (Makel & Plucker, 2014). This discrepancy may be due to sampling bias, socially desirable responding, or education researchers engaging in research that is replicative but not explicitly identified as a replication (see Cook et al., 2016). Across all Open Practices, participating education researchers self-reported a higher level of engagement than they expected from others in the field. This may be due to sampling bias, socially desirable responding, or lack of awareness of other researchers' engagement in Open Practices. Although few respondents indicated Open Practices should never be used, a small proportion felt this way for each specific practice. Notably, 5% of education researchers indicated that research should never be made openly accessible. Our future qualitative analysis will explore the reasons participants provided for why they feel it is appropriate to engage in QRPs, and why a small percentage feel it is never appropriate to engage in open practices.

The large percentage of education researchers admitting to using QRPs brings into question whether published education research findings are reliable. Moreover, such self-admission of QRP use may help explain publication bias (e.g., Polanin et al., 2016; Sham & Smith, 2014) and the scarcity of negative findings (Fanelli, 2011; Pigott et al., 2013; Polanin et

al., 2016) in the education research base, as well as why many education interventions fail to scale up (Elmore, 1996; Lortie-Forgues & Inglis, 2019). These may all be consequences of study effects being inflated by the application of QRPs.

Who Uses QRPs and Open Practices?

Our results appear largely aligned with those from psychology, indicating that QRP use is not rare in education research. With a few exceptions, use of QRPs does not appear to vary by geographic location, career stage, or quantitative orientation. This suggests that the use of QRPs is prevalent across the field and not within some small pockets of the community.

Our results suggest that Open Practices are not rare in education research and that geographic location or quantitative orientation has little bearing on their use. Early career researchers reported lower rates of Open Practice use than their more experienced colleagues. That may be due having had fewer opportunities to do so or out of concern that they would not be appropriate. Further research is needed to investigate.

Future Directions

If the field wishes to develop consensus over what research behaviors are acceptable and decrease use of QRPs, several steps could be taken by various stakeholders. One broad route to decrease QRPs (or make their use more transparent) is to increase the use of Open Practices by researchers. Most Open Practices help avoid QRPs in new research and conducting replications can help identify accumulated QRP debt that has already accrued in the existing literature.

Additionally, Makel and colleagues (2019) proposed several large-scale collaboration options that education researchers can use to identify which results in the existing literature can be replicated. Journal publishers and editors in education might encourage and reward Open Practices by adopting the Transparency and Openness Promotion (TOP) guidelines (Nosek et al.,

2015), recognizing Open Practices by awarding electronic badges (Kidwell et al., 2016), allowing for submissions of Registered Reports (Chambers, 2019), and requiring that researchers disclose methodological and design specifics (LeBel et al., 2013). Other ways stakeholders can facilitate Open Practices include reviewers asking for data and replicating reported analyses, professional societies adopting the TOP guidelines, and funders mandating that funded researchers use Open Practices.

Future work should seek to replicate our estimates (or similar estimates) to inform the field more clearly where it stands regarding questionable and open research practices. For example, future work could more directly estimate how many (and which) results are based on QRPs by using publication as the unit of analysis instead of researcher. Future work could also explore the extent to which QRP use is driven by intentional subversion of accepted norms vs. ignorance of those norms. Additional work could also try to compare beliefs about appropriateness with self-reported use of that behavior.

More broadly, future work can help shed light on the state of the field's research. For example, does appropriateness of some behaviors vary across quantitative and qualitative methods? Are there behaviors that are more appropriate in education research compared to other fields? Are they appropriate in some contexts/countries but not others? All these questions require greater understanding around the use of QRPs and Open Practices. Further work that helps create and entrench cultural norms in educational research for what constitutes appropriate research behavior will help create benchmarks for how well the field's behaviors align with research consumer (policymakers, practitioners, and other researchers) beliefs and aspirations for the field.

Conclusion

Our results suggest QRPs are used by many education researchers. Like airbrushing a photograph, QRPs may smooth unsightly blemishes, creating a false sense of desirable appearance. However, QRPs also have the potential to distort research results to the point of unrecognizability, making them unrealistically attractive. Regardless of where on the continuum of severity QRPs fall, the consequences of creating a research literature littered with such practices are real.

At the same time, our results provide evidence that education researchers generally disagreed on the appropriateness of both questionable and open research practices, suggesting there is not yet consensus on which research behaviors are misconduct and which are accepted practices. The field needs to increase consensus over which practices are misconduct and which research behaviors are recommended in education research.

QRPs can lead to massive misinvestment in educational practices with unrealistic expectations of impact. With public school expenditures in the U.S. exceeding \$730 billion (Hussar et al., 2020), the financial stakes are enormous. More importantly, implementing educational practices based on faulty research can cause real harm to children and distract from implementing practices based on less-alluring but more reliable and accurate research. Greater implementation of Open Practices can help change norms in education research while improving both research transparency and research quality. We recommend the education research community develop infrastructure, provide incentives, and establish policies encouraging the use of Open Practices. Doing so can help to increase the credibility, effectiveness, and relevance of education research.

References

- Agnoli, F., Wicherts, J. M., Veldkamp, C. L. S., Albiero, P., & Cubelli, R. (2017). Questionable research practices among Italian research psychologists, *PloS One*, *12*(3), e0172792 doi: 10.1371/journal.pone.0172792
- Allen, C., & Mehler, D. M. (2019). Open Science challenges, benefits and tips in early career and beyond. *PLoS Biology*, *17*(5), e3000246. doi: 10.1371/journal.pbio.3000246
- Begley, C. G., & Ellis, L. M. (2012). Raise standards for preclinical cancer research. *Nature*, *483*, 531–533. doi: 10.1038/483531a
- Campbell, H. A., Micheli-Campbell, M. A., & Udyawer, V. (2019). Early career researchers embrace data sharing. *Trends in Ecology & Evolution*, *34*, 95-98. doi: 10.1016/j.tree.2018.11.010
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T. H., Huber, J., Johannesson, M., ... & Altmejd, A. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*, *2*, 637-644. doi: 10.1038/s41562-018-0399-z
- Chambers, C. (2019). The registered reports revolution Lessons in cultural reform. *Significance*, *16*(4), 23-27. doi: [10.1111/j.1740-9713.2019.01299.x](https://doi.org/10.1111/j.1740-9713.2019.01299.x)
- Chambers, C. D., Feredoes, E., Muthukumaraswamy, S. D., & Etchells, P. J. (2014). Instead of “playing the game” it is time to change the rules: Registered reports at AIMS Neuroscience and beyond. *AIMS Neuroscience*, *1*, 4-17. doi:10.3934/Neuroscience2014.1.4

- Choshen-Hillel, S., Shaw, A., & Caruso, E. M. (2020). Lying to appear honest. *Journal of Experimental Psychology: General*, 149(9), 1719–1735.
<https://doi.org/10.1037/xge0000737>
- Cook, B. G., Collins, L. W., Cook, S. C., & Cook, L. H. (2016). A replication by any other name: A systematic review of replicative intervention studies. *Remedial and Special Education*, 37, 223-234. doi: 10.1177/0741932516637198
- Cook, B. G., Lloyd, J. W., Mellor, D., Nosek, B. A., & Therrien, W. J. (2018). Promoting open science to increase the trustworthiness of evidence in special education. *Exceptional Children*, 85, 104-118. doi: 10.1177/0014402918793138
- Elmore, R. (1996). Getting to scale with good educational practice. *Harvard Educational Review*, 66, 1-27. doi: 10.17763/haer.66.1.g73266758j348t33
- Fanelli, D. (2012). Negative results are disappearing from most disciplines and countries. *Scientometrics*, 90, 891–904. doi: 10.1007/s11192-011-0494-7
- Fanelli, D., Costas, R., & Larivière, V. (2015). Misconduct policies, academic culture and career stage, not gender or pressures to publish, affect scientific integrity. *PloS One*, 10(6), e0127556. doi: 10.1371/journal.pone.0127556
- Fanelli, D., & Ioannidis, J. P. (2013). US studies may overestimate effect sizes in softer research. *Proceedings of the National Academy of Sciences*, 110, 15031-15036. doi: 10.1073/pnas.1302997110
- Farnham, A., Kurz, C., Öztürk, M. A., Solbiati, M., Myllyntaus, O., Meekes, J., ... & Kanninen, L. (2017). Early career researchers want Open Science. *Genome Biology*, 18(1), 221. doi: 10.1186/s13059-017-1351-7

- Fiedler, K., & Schwarz, N. (2015). Questionable research practices revisited. *Social Psychological & Personality Science*, 7, 45–52. doi: 10.1177/1948550615612150
- Fraser, H., Parker, T., Nakagawa, S., Barnett, A., & Fidler, F. (2018). Questionable research practices in ecology and evolution. *PLoS ONE*, 13(7), e0200303. doi: 10.1371/journal.pone.0200303
- Gehlbach, H. & Robinson, C. D. (2018). Mitigating illusory results through preregistration in education. *Journal of Research on Educational Effectiveness*, 11, 296-315. doi: 10.1080/19345747.2017.1387950
- Halevy, R., Shalvi, S., & Verschuere, B. (2014). Being honest about dishonesty: Correlating self-reports and actual lying. *Human Communication Research*, 40(1), 54-72. <https://doi.org/10.1111/hcre.12019>
- Horowitz, J. L. (2001). The bootstrap. In *Handbook of econometrics* (Vol. 5, pp. 3159-3228). Elsevier.
- Hussar, B., Zhang, J., Hein, S., Wang, K., Roberts, A., Cui, J., Smith, M., Bullock Mann, F., Barmer, A., & Dilig, R. (2020). The Condition of Education 2020 (NCES 2020-144). U.S. Department of Education. Washington, DC: National Center for Education Statistics. Retrieved from: <https://nces.ed.gov/pubsearch/pubsinfo.asp?pubid=2020144>.
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23, 524–532. doi: 10.1177/0956797611430953
- Kaplan RM, Irvin VL (2015) Likelihood of null effects of large NHLBI clinical trials has increased over time. *PLoS ONE*, 10(8), e0132382. doi: 10.1371/journal.pone.0132382

- Kerr, N. L. (1998). HARKing: Hypothesizing after the results are known. *Personality and Social Psychology Review*, 2, 196–217. Doi: 10.1207/s15327957pspr0203_4
- Kidwell, M. C., Lazarević, L. B., Baranski, E., Hardwicke, T. E., Piechowski, S., Falkenberg, L. S., ... & Errington, T. M. (2016). Badges to acknowledge open practices: A simple, low-cost, effective method for increasing transparency. *PLoS Biology*, 14(5), e1002456. doi: 10.1371/journal.pbio.1002456
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Jr., Alper, S., ... Nosek, B. A. (2018). Many labs 2: Investigating variation in replicability across sample and setting. *Advances in Methods and Practices in Psychological Science*, 1, 443-490. doi: 10.1177/2515245918810225
- Lakens, D. (2017). Equivalence tests: A practical primer for t-tests, correlations, and meta-analyses. *Social Psychological and Personality Science*, 8, 355-363. doi: 10.1177/1948550617697177
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: A practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, 4, 863. doi: 10.3389/fpsyg.2013.00863
- LeBel, E. P., Borsboom, D., Giner-Sorolla, R., Hasselman, F., Peters, K. R., Ratliff, K. A., & Smith, C. T. (2013). PsychDisclosure.org: Grassroots support for reforming reporting standards in psychology. *Perspectives on Psychological Science*, 8, 424-432. doi: [10.1177/1745691613491437](https://doi.org/10.1177/1745691613491437)
- Lortie-Forgues, H. & Ingles, M., (2019). Most rigorous large-scale educational RCTs are uninformative: Should we be concerned? *Educational Researcher*, 48, 158-166. doi: 10.3102/0013189X19832850

Makel, M. C. & Plucker, J. A. (2014). Facts are more important than novelty: Replication in the education sciences. *Educational Researcher*, 43, 304-316. doi:

[10.3102/0013189X14545513](https://doi.org/10.3102/0013189X14545513)

Makel, M. C., Smith, K. N., McBee, M. T., Peters, S. J., & Miller, E. M. (2019). A path to greater credibility: Large-scale collaborative education research. *AERA Open*, 5(4).

<https://journals.sagepub.com/doi/full/10.1177/2332858419891963>

Nosek, B. A., Alter, G., Banks, G. C., Borsboom, D., Bowman, S. D., Breckler, S. J., ... &

Yarkoni, T. (2015). Promoting an open research culture. *Science*, 348(6242), 1422- 1425.

doi: 10.1126/science.aab2374

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science.

Science, 349(6251), aac4716–aac4716. doi: 10.1126/science.aac4716

Pashler, H., & Harris, C. R. (2012). Is the replicability crisis overblown? Three arguments examined. *Perspectives on Psychological Science*, 7, 531–536. doi:

10.1177/1745691612463401

Pigott, T. D., Valentine, J. C., Polanin, J. R., Williams, R. T., & Canada, D. D. (2013). Outcome-reporting bias in education research. *Educational Researcher*, 42, 424–432. doi:

[10.3102/0013189X13507104](https://doi.org/10.3102/0013189X13507104)

Polanin, J. R., Tanner-Smith, E. E., & Hennessy, E. A. (2016). Estimating the difference between published and unpublished effect sizes: A meta-review. *Review of Educational Research*,

86, 207-236. doi: 10.3102/0034654315582067

Prinz, F., Schlange, T., & Asadullah, K. (2011). Believe it or not: How much can we rely on

published data on potential drug targets? *Nature Reviews Drug Discovery*, 10, 712–713.

- Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological Bulletin*, 86, 638–641. doi: 10.1037//0033-2909.86.3.638
- Rotton, J., Foos, P. W., Vanmeek, L., & Levitt, M. (1995). Publication practices and the file drawer problem: A survey of published authors. *Journal of Social Behavior and Personality*, 10, 1–13.
- Sham, E., & Smith, T. (2014). Publication bias in studies of an applied behavior-analytic intervention: An initial analysis. *Journal of Applied Behavior Analysis*, 47, 663–678. doi: 10.1002/jaba.146
- Simmons, J., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359–1366. doi: [10.1177/0956797611417632](https://doi.org/10.1177/0956797611417632)
- Simmons, J. P., Nelson, L. D., Simonsohn, U. (2018). False-positive citations. *Perspectives on Psychological Science*, 13(2), 255-259. <https://doi.org/10.1177/1745691617698146>
- Spellman, B. A. (2012). Introduction to the special section: Data, data, everywhere . . . especially in my file drawer. *Perspectives on Psychological Science*, 7, 58–59. doi: 10.1177/1745691611432124
- van der Zee, T. & Reich, J. (2018). Open education science. *AERA Open*, 4, 1-15. doi: 10.1177/2332858418787466
- Vazire, S. (2018). Implications of the credibility revolution for productivity, creativity, and progress. *Perspectives on Psychological Science*, 13, 411-417. doi: 10.1177/1745691617751884

Table 1

Descriptive statistics for prevalence estimates, self-reported use, and appropriateness estimates for questionable and open research practices

Item Stem	Abbreviation	Estimated percentage of colleagues who engage in the practice (Number of responses)	Percentage indicating they have engaged in the practice at least one time (Number of responses)	Percentage indicating the practice should never be used (Number of responses)
Questionable Research Practices				
1. Not reporting studies or variables that failed to reach statistical significance (e.g., $p \leq .05$) or some other desired statistical threshold.	Omitting nonsignificant studies or variables	60.75% (764)	61.69% (783)	37.84% (761)
2. Not reporting covariates that failed to reach statistical significance (e.g., $p \leq .05$) or some other desired statistical threshold.	Omitting nonsignificant covariates	49.51% (752)	41.64% (773)	44.67% (750)
3. Reporting an unexpected finding or a result from exploratory analysis (i.e., not explicitly planned in advance) as having been predicted from the start. ¹	HARKing	47.26% (853)	45.80% (880)	55.09% (864)
4. Reporting a set of results as the complete set of analyses when other analyses were also conducted.	Omitting analyses	60.21% (849)	67.28% (871)	26.49% (838)
5. Rounding-off a p value or other quantity to meet a pre-specified threshold (e.g., reporting $p = .054$ as $p = .05$ or $p = .013$ as $p = .01$).	Rounding p-values	41.95% (785)	28.61% (819)	67.83% (802)

6. Deciding to exclude data points after first checking the impact on statistical significance (e.g., $p \leq .05$) or some other desired statistical threshold.	Data Exclusion ARKing	36.67% (750)	25.48% (806)	65.63% (768)
7. Collecting more data for a study after first inspecting whether the results are statistically significant (e.g., $p \leq .05$).	Data peeking	36.38% (773)	28.66% (806)	37.52% (781)
8. Changing to another type of statistical analysis after the analysis initially chosen failed to reach statistical significance (e.g., $p \leq .05$) or some other desired statistical threshold.	Analysis Gaming	49.04% (789)	49.57% (811)	34.47% (789)
9. Not disclosing known problems in the method, analysis, or data quality that potentially impact conclusions.	Hiding methodological problems	41.70% (861)	24.07% (889)	80.97% (883)
10. Filling in missing data points without identifying those data as simulated.	Filling in missing data	23.46% (847)	9.69% (898)	82.39% (886)
Open Practices				
11. Preregistering hypotheses and analysis plans prior to data collection.	Preregistration	30.16% (849)	53.92% (873)	3.94% (837)
12. Sharing data you collected to a publicly accessible, online repository.	Data Sharing	24.78% (858)	45.61% (888)	3.70% (866)
13. Sharing code or other research materials to a publicly accessible, online repository.	Materials Sharing	29.45% (853)	58.94% (884)	2.32% (862)

14. Sought to replicate the work of other researchers by following their methods as closely as possible with no intentional changes.	Replication	26.65% (848)	43.26% (876)	2.01% (845)
15. Posted copies of your research so that it is not behind a paywall (e.g., on a publicly accessible, online preprint server).	Open Access	50.23% (861)	78.43% (881)	5.03% (855)

Note. ¹The wording of this item has an additional change from Fraser et al (2018). We removed the word “statistical” so that the item would be applicable to both quantitative and qualitative research.

Figure 1. Prevalence estimates of questionable and open research practices. Note. Error bars represent the standard deviation of the prevalence estimate percentage. *Note. ONSV stands for Omitting Nonsignificant Studies or Variables, ARKing for after results are known, and HARKing for hypothesizing after results are known.

Figure 2. Self-reported use of questionable and open research practices. *Note. ONSV stands for Omitting Nonsignificant Studies or Variables, ARKing for after results are known, and HARKing for hypothesizing after results are known.

Figure 3. Appropriateness estimates of questionable and open research practices. *Note. ONSV stands for Omitting Nonsignificant Studies or Variables, ARKing for after results are known, and HARKing for hypothesizing after results are known.